On Visibility Representations of Non-planar Graphs*

Therese Biedl[1], Giuseppe Liotta[2], Fabrizio Montecchiani[2]

David R. Cheriton School of Computer Science, University of Waterloo, Canada
biedl@uwaterloo.ca

[2]Dipartimento di Ingegneria, Università degli Studi di Perugia, Italy
{giuseppe.liotta,fabrizio.montecchiani}@unipg.it

arXiv:1511.08592v1 [cs.CG] 27 Nov 2015

Keywords: Visibility Representations, 1-Planarity, Recognition Algorithm, Forbidden Configurationor

Abstract

A rectangle visibility representation (RVR) of a graph consists of an assignment of axis-aligned rectangles to vertices such that for every edge there exists a horizontal or vertical line of sight between the rectangles assigned to its endpoints. Testing whether a graph has an RVR is known to be NP-hard. In this paper, we study the problem of finding an RVR under the assumption that an embedding in the plane of the input graph is fixed and we are looking for an RVR that reflects this embedding. We show that in this case the problem can be solved in polynomial time for general embedded graphs and in linear time for 1-plane graphs (i.e., embedded graphs having at most one crossing per edge). The linear time algorithm uses a precise list of forbidden configurations, which extends the set known for straight-line drawings of 1-plane graphs. These forbidden configurations can be tested for in linear time, and so in linear time we can test whether a 1-plane graph has an RVR and either compute such a representation or report a negative witness. Finally, we discuss some extensions of our study to the case when the embedding is not fixed but the RVR can have at most one crossing per edge.

1 Introduction

A visibility representation is an appealing method of displaying graphs. It consists of assigning axis-aligned rectangles to vertices and horizontal or vertical line segments to edges in such a way that line segments of edges begin and end at their endpoints and intersect no rectangles in-between. Equivalently, each edge corresponds to a horizontal or vertical line-of-sight between its endpoints; hence the name "visibility". Edge segments may cross each other, but any such crossing occurs at a right angle. These right-angle crossings as well as the absence of bends leads to a high readability of the visualization. See Fig. 1 for an illustration.

Visibility representations were introduced to the Computational Geometry community in 1985, when Wismath [32] showed that every planar graph has a bar-visibility representation, i.e., where vertices are represented as horizontal bars and edges correspond to vertical visibilities. (The same result was discovered independently multiple times [12, 22, 24, 29, 30].)

*Research of Therese Biedl supported by NSERC. Research of Giuseppe Liotta and Fabrizio Montecchiani supported in part by the MIUR project AMANDA "Algorithmics for MAssive and Networked DAta", prot. 2012C4E3KT_001. Research undertaken while Fabrizio Montecchiani was visiting the University of Waterloo supported by NSERC.

Such bar-visibility representations can exist only for planar graphs, so for representing non-planar graphs one must use edge segments in both directions. (We use the term *rectangle visibility representation* or *RVR* to distinguish such drawings from bar-visibility representations.) Clearly not any graph G can have an RVR either; for example G must have thickness 2 (i.e., be the union of two planar graphs), and, as was shown in [3], it must have at most $6n - 20$ edges. But neither of these conditions is sufficient. In fact, Shermer showed in 1996 that it is NP-hard to test whether a given graph has an RVR [25].

In this paper, we study the problem of testing whether a graph has a rectangle visibility representation if some aspects of it are fixed a priori. In previous work by Streinu and Whitesides [26], it was shown that testing whether a graph has an RVR is polynomial if the following additional information is given: we know for each edge whether it is drawn horizontal or vertical, we know the rotational scheme (i.e., the fixed order of edges around each vertex), and we know which vertices form the outer face of the visibility representation. Streinu and Whitesides give necessary and sufficient conditions for when such restrictions can be realized, and show that these can be tested in $O(n^2)$ time.

In this paper, we study a slightly different scenario. Like Streinu and Whitesides, we assume that the rotational scheme and the outer face is fixed. However, we do not require to know the directions of the edges; instead we must know a priori which edges cross each other, the order in which these crossings occur along each edge (if there is more than one crossing on an edge). In other words, the graph has a *fixed embedding* (see also Section 2). It turns out that a reduction to the topology-shape-metrics approach then allows to test in $O(n^{1.5} \log n)$ time whether a graph with n vertices has a rectangle visibility representation; we give the details in Section 2.3.

As our main result, we then consider the case when any edge contains at most one crossing (the so-called *1-planar graphs*). Previous attempts to create visibility representations for 1-planar graphs had relaxed the restriction that edges must not cross non-incident vertices, and defined a *bar k-visibility drawing* to be a bar-visibility drawing where each line of sight can intersect at most k bars [9]. In other words, each edge can cross at most k vertices. These crossings are arguably less intuitive than crossings between edges. Brandenburg *et al.* [4] and independently Evans *et al.* [14] prove that 1-planar graphs have a bar 1-visibility drawing (if we are allowed to change the embedding).

In contrast to these result, we consider in Section 3 embedding-preserving rectangle visibility representations of embedded 1-planar graphs, also called 1-plane graphs. It is easy to see that not all 1-plane graphs have such an RVR (see also Section 4). The main contribution of our paper is to show that for a 1-plane graph it is possible to test in linear time whether it has an RVR. The approach is entirely different from the topology-shape-metrics above and reveals much insight into the structure required for an RVR to exist. Specifically, we consider three configurations (called B-configuration, W-configuration, and the newly defined T-configuration) and show that a 1-plane graph has an RVR if and only if none of these three configurations occurs. This mirrors in a pleasing way the result by Thomassen [31], which shows that a 1-plane graph has a straight-line drawing if and only if it contains no B-configuration or W-configuration. Our result also provides an answer (in a special case) to Shermer's question of characterizing when graphs have RVRs [25]. Also, we prove that testing whether a 1-plane graph contains any of the three forbidden configurations can be done in linear time, and in the absence of them we can compute in linear time an RVR which fits into a grid of quadratic size. It may be worth recalling that embedding-preserving straight-line drawings of 1-plane graphs may require exponential area [17].

In contrast to planar graphs, 1-planar graphs may have an exponential number of different embeddings even if 3-connected or 4-connected (see, e.g., Fig. 9). Hence our algorithm does not automatically cover the case when the embedding is not fixed. We briefly study 1-planar graphs without fixed embedding in Section 4. We show that there exist infinitely many 3-connected

1-planar graphs not admitting a 1-planar embedding realizable as an RVR, even after removing linearly many edges. On the positive side, we show that 4-connected 1-planar graphs, which include the optimal 1-planar graphs, always admit an RVR, except for at most one edge which cannot be represented.

Paper structure. The remainder of this paper is organized as follows. In Section 2 we introduce basic definitions and results which will be used throughout the paper. In Section 3 we characterize the 1-plane graphs that admit an RVR. In Section 4 we show both negative and positive results in the variable embedding setting. Finally, in Section 5 we conclude with a discussion and open problems.

2 Definitions and Basic Results

2.1 Embedding-preserving Visibility Representations

We only consider *simple graphs*, i.e., graphs with neither loops nor multiple edges. A *drawing* Γ of a graph G maps each vertex to a point of the plane and each edge to a Jordan arc between its two endpoints. We only consider *simple drawings*, i.e., drawings such that the arcs representing two edges have at most one point in common, which is either a common endpoint or a common interior point where the two arcs properly cross. A drawing is *planar* if no two arcs cross. A drawing divides the plane into topologically connected regions, called *faces*. The unbounded region is called the *outer face*. For a planar drawing the boundary of a face consists of vertices and edges, while for a non-planar drawing the boundary of a face may contain vertices, crossings, and edges (or parts of edges). An *inner* face/edge/vertex is a face/edge/vertex that is not (part of) the outer face.

A *rotational scheme* of a graph G consists of assigning at each vertex a cyclic order of the edges incident to that vertex. A drawing is said to respect a rotational scheme if scanning around the vertex in clockwise order encounters the edges in the prescribed order. We say that a graph has a *fixed embedding* if the graph comes with a rotational scheme, a fixed order of crossings along each edge, and with one face indicated to be the outer face. (One can show that fixing the rotational scheme and the order of crossings fixes all faces.) Given a graph G with a fixed embedding, the *planarization* G_p of G is the graph obtained by replacing each crossing point of G with a *dummy-vertex*. The vertices of G in G_p are called *original vertices* to avoid confusion. The planarization G_p is a *plane* graph, i.e., a planar graph with a fixed embedding without crossing.

A *visibility representation* (also called *rectangle visibility representation* or RVR for short) of a graph G is an assignment of disjoint axis-aligned rectangles to the vertices in such a way that for every edge there exists a horizontal or vertical line of sight between its endpoints. We follow a commonly adopted model where the lines of sight are *thick*, i.e., have non-zero area (see e.g. [20, 26, 29, 32]). Thus we may assume that each edge is routed as a (horizontal or vertical) line segment, attaching at each end at a point that is not a corner of the rectangle of that vertex. In what follows, when this leads to no confusion, we shall use the same term *edge* to indicate both an edge of a graph, the line of sight between two rectangles, and the line segment representing the edge, and we use the same term *vertex* for both the vertex of a graph and the rectangle that represents it. We observe that if we have an RVR, then we can naturally extract a drawing from it as follows. Place a point for each vertex v inside the rectangle representing v and connect it to all the attachment points of incident edges of v inside the rectangle; this adds no crossing. We say that an RVR Γ of a graph G *respects* the embedding of G if applying this operation to Γ results in the same embedding.

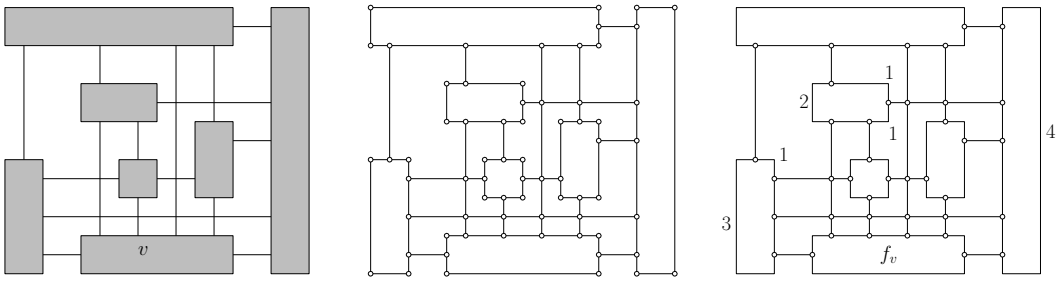

Figure 1: A visibility representation of a graph, the orthogonal representation obtained from it, and the orthogonal representation used to compute it. We show only selected non-zero bend-counts.

2.2 Topology-shape-metrics

The *topology-shape-metrics* approach is a well-known method introduced by Tamassia [28] to create orthogonal drawings of graphs, i.e., drawings where all vertices are points and edges are curves with horizontal and vertical segments. In this section, we briefly review how any RVR can be expressed as an orthogonal-shape-metric; this is quite straightforward but has to our knowledge not been used before. We will use such metrics twice in our paper: once to determine whether a given embedded graph has an RVR, and once to store an RVR obtained for 4-connected 1-plane graphs so that they can be manipulated efficiently.

Consider an RVR $\mathcal{R}$ and interpret it as a PSLG, i.e., a planar straight-line graph in the computational geometry sense. Thus, replace every corner of a rectangle, every attachment point of an edge at a rectangle, and every crossing by a node of the PSLG, and connect nodes with arcs if and only if they are connected by (part of) an edge segment or (part of) the boundary of a rectangle. Each node of the PSLG has around it a sequence of angles that are all multiple of $90°$ and sum to $360°$. Likewise each face of the PSLG has around it a sequence of angles that are all multiple of $90°$ and that sum to $(k-2)180°$ for inner faces and $(k+2)180°$ for the outer face, where k is the number of nodes on the face. This motivates the following definition:

Definition 1. *An* orthogonal representation *is a plane graph H with maximum degree 4, an assignment of an angle $\alpha(v,f) \in \{90°, 180°, 270°\}$ to any vertex-face-incidence (between vertex v and face f) of H, and a bend-count $b(e,f) \in \{0, +1, -1, +2, -2, \ldots\}$ to any edge-face-incidence (between edge e and face f) of H, such that the following is satisfied:*

- *at every vertex v, the angles at v sum to $360°$,*
- *at every edge e, the two bend-counts sum to 0,*
- *for every inner face f with k vertices, the angles and bend-counts at f sum to $(k-2)180°$. Formally,*
$$\sum_{(v,f)} \alpha(v,f) + \sum_{(e,f)} b(e,f) \cdot 90° = (k-2)180°$$
- *the angles and bend-counts at the outer face sum to $(k+2)180°$, where k is the number of vertices on the outer face.*

We have shown above:

Observation 1. *If G has an RVR Γ, then construct a plane graph H by replacing every crossing, every edge-vertex attachment point and every corner of a vertex-rectangle of Γ with a node in H. The resulting graph then has an orthogonal representation where all bend-counts are 0.*

Tamassia showed [28] that for any orthogonal representation, one can reconstruct a planar orthogonal drawing of the graph that respects these angles and bend-counts. In other words, at any vertex-face-incidence the angle is $\alpha(v, f)$, and any edge e has $|b(e, f)|$ bends, where the bends have angle 90° in the face where the bend-count at e is positive and hence angle $-90^\circ = 270^\circ$ in the other face. In particular, it is not required to know the edge-lengths a priori. This reconstruction process takes linear time in the size of H and the total number of bend-counts. If the orthogonal representation came from an RVR, then the resulting drawing has the exact same set of crossings and the face that came from a vertex rectangle is again rectangular, so the resulting drawing is again an RVR.

Tamassia also showed how to compute for a given plane graph H some angle-assignment and bend-counts such that the result is an orthogonal representation. This is done via a flow-approach; see [28] or [10] for details. If one uses a minimum-cost-flow algorithm, then this allows to compute the topological orthogonal shape that minimizes the total number of bends. Of higher interest to us is instead the following variant: One can specify upper and lower bounds on the angles and bend-counts, and can then use a maximum-flow algorithm to test whether a feasible flow exists, and hence to test whether there exists an orthogonal representation for which angles and bend-counts are within the specified bound. The first version can be tested in $O(n^{1.5})$ time while the second version in $O(n^{1.5} \log n)$ time, where n is the number of vertices of H [7].

2.3 Testing embedded graphs for RVRs

Now presume that we are given a graph G with a fixed embedding. In this section, we will show how to test whether G has an RVR that respects its embedding using the topology-shape-metrics approach. Note that the embedding tells us nearly the entire graph H of the topology of a (putative) RVR of G. The only thing unknown is the location of the corners of each rectangle of a vertex. We can now interpret these corners as bends (rather than as nodes) in the topology. Hence define a graph H as follows (see also Fig. 1):

- Start with the planarization G_p of G.
- For any original vertex v of G, define a cycle C_v with $\deg(v)$ nodes in H. Attach the (parts of) edges incident to v to the nodes of cycle C_v in the specified order of the planar embedding, in such a way that the interior of the cycle forms a face. In other words, form a cycle that can represent the boundary of the rectangle of v.
- Impose an upper and lower bound of 0 onto the bend-counts of all (parts of) edges of G. In other words, we do not allow any bends for the original edges.
- Impose a lower bound of 0 onto each bend-count $b(e, f_v)$ where e is an edge of some cycle C_v and f_v is the face formed by cycle C_v. In other words, cycle C_v may have bends, but all such bends must form 90° angles towards f_v.
- If a is a node of C_v, then impose an upper and lower bound of 180° onto the angle $\alpha(a, f_v)$, where f_v as before is the face inside C_v. In other words, at attachment points of edges face f_v must not have a corner.

Theorem 1. *Let G be an n-vertex graph with a fixed embedding. Then we can test in $O(n^{1.5} \log n)$ time whether G has an RVR that respects this embedding.*

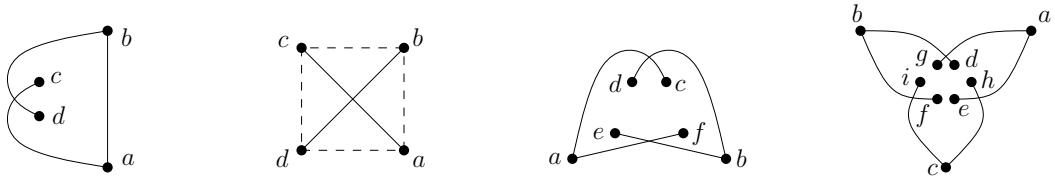

(a) B (b) Kite (c) W (d) T

Figure 2: Possible crossing configurations in a 1-plane graph G. (a) A B-configuration. (b) A (subgraph) of a kite. The dashed edges either do not exist or define faces. (c) A W-configuration. (d) A T-configuration.

Proof. Build the graph H as described above. It should be obvious that if we can find an orthogonal representation of H that respects these constraints, then it gives rise to an RVR of G. Indeed, for any vertex v the cycle C_v must form a rectangle: It cannot have a bend at any attachment-point, and any bend that it has provides an angle of 90°, hence by the condition it has exactly four such bends. At any dummy-vertex the four angles must be 90° by the condition for nodes. All edges of H that came from edges of G are drawn straight, and continue straight at any crossing, so give rise to a line of sight. We then use Tamassia's flow-approach to test whether an orthogonal representation of H exists that satisfies all constraints. This takes $O(n^{1.5} \log n)$ time [7] (G has $O(n)$ crossings). □

3 Embedded 1-planar graphs

While the topology-shape-metrics approach gives a polynomial way to test whether a graph has an RVR, its black-box approach is somewhat unsatisfactory. In particular, in case of a negative answer, we obtain no good insights as to which parts of the graph prevented the existence of a representation. In this section, we therefore turn our attention to 1-planar graphs. We chose this graph class for two reasons. One is that they are known to have thickness 2 and at most $4n - 8$ edges [23], and so they are good candidates for always having an RVR. (We will, however, see that this is not the case.) Secondly, 1-planar graphs have been studied widely in the graph theory and graph drawing communities (see e.g. [1, 4, 8, 14, 17, 21, 23, 27, 31]), and visibility representations of 1-planar graphs hence should be of interest.

3.1 Background

A graph is *1-planar* if it has a drawing with at most one crossing per edge. A *1-plane* graph is a 1-planar graph with a fixed embedding. Recall that this means that we have a fixed order of edges at each vertex, it is fixed which pairs of edges cross, and we know what the faces are and which one of them is the outer face. A 1-plane graph G is called *triangulated 1-plane* if all faces are triangles. In a 1-plane graph, a facial triangle is composed of three vertices or two vertices and a crossing point.

Let $e = (a, c)$ and $e' = (b, d)$ be two edges of G that cross at a point p. We say that e, e' induce a *B-configuration* if there exists an edge between their endpoints (say edge (a, b)) such that the triangle $\{a, b, p\}$ contains vertices c and d inside (see also Fig. 2a).

If the crossing is not in a B-configuration, then there are two further possibilities for edge (a, b). The triangle $(a, b), (b, p), (p, a)$ may be an interior face, or it could have other vertices inside. , even in a 3-connected graph since edge (a, b) may be crossed. If, for all pairs in $\{a, c\} \times \{b, e\}$,

there either is no edge between the vertices or the triangle that it forms is a face, then we could add the edges that did not exist previously, routing them along the crossing, and then obtain a *kite*, i.e., a crossing with a 4-cycle among its endpoints such that removing the crossing turns the 4-cycle into a face.

Let (a, g) and (b, d) be two edges of G that cross at a point p, and let (a, f) and (b, e) be two further edges of G that cross at a point q. The four edges induce a *W-configuration* if vertices c, d, e, f lie inside the closed region delimited by the edge segments (a, p), (b, p), (a, q), and (b, q) (see also Fig. 2c).

For the result in this paper, we need to introduce a forth configuration, called the *trillium configuration*, or *T-configuration* for short, which is illustrated in Fig. 2d. Namely, let (a, c) and (b, d) be a pair of edges of G that cross at a point p. Let (a, e) and (c, h) be a second pair of edges that cross at a point q. Let (b, f) and (c, i) be a third pair of edges that cross at a point t. The six edges induce a T-configuration if vertices d, e, f, g, h, i lie inside the closed region delimited by the edge segments (a, p), (b, p), (a, q), (c, q), (b, t), and (c, t). Vertices a, b, c are called the *outer vertices* of the T-configuration, while the remaining six vertices are the *inner vertices*.

3.2 Main result and outline

In this section we give a characterization of 1-plane graphs that admit an RVR. Our characterization is summarized by the following theorem.

Theorem 2. *A 1-plane graph G admits an RVR if and only if it contains no B-configuration, no W-configuration, and no T-configuration.*

Notably, Theorem 2 extends the characterization for straight-line drawability of 1-plane graphs given by Thomassen [30], adding the T-configuration to the set of obstructions. The necessity of the condition is proved via an easy angle-counting-argument in Section 3.3. The sufficiency-proof is quite intricate and is given in Section 3.4. It comes with an efficient algorithm to construct an RVR Γ of a drawable graph G. Finally, we describe a testing procedure (Section 3.6) that easily follows from our characterization. Next theorem is obtained by combining our testing procedure and our drawing algorithm.

Theorem 3. *Let G be a 1-plane graph with n vertices. There exists an $O(n)$-time algorithm to test whether G admits an RVR. Also, in the positive case the algorithm computes an RVR of G that has $O(n^2)$ area.*

3.3 Proof of necessary condition

Our first goal is to show that no B-/W-/T-configuration can be a subgraph of a 1-plane graph admitting an RVR. To this aim we show a slightly stronger statement.

Lemma 1. *Let G be a graph that admits an RVR Γ and whose outer face is a cycle $\mathcal{C}$. If while walking along $\mathcal{C}$ we encounter k vertices and c crossing points, then $k \geq 3$ and $c \leq 2k - 4$.*

Proof. Convert Γ into an orthogonal representation (Observation 1) and consider the angles at the outer face. Since there are k vertices on the outer face, we have $2k$ edge-attachment points, and each contributes $90°$. Each of the c crossings contributes $90°$. If b is the number of rectangle-corners on the outer face, then $b \leq 4k$, and each contributes $270°$. In total the outer face has $2k + c + b$ nodes and the sum of angles must be $180°(2k + c + b + 2)$. But we also know that the angles sum to $90°(2k + c) + 270°b = 180°(2k + c + b + 2) + 90°(b - 2k - c - 4)$. So $b - 2k - c - 4 = 0$, hence $4k \geq b = 2k - c - 4$, hence $c \leq 2k - 4$. This implies $k \geq 2$, and in fact $k \geq 3$ is required, else $c = 2$ and the outer face would consist of a double edge. □

Since a B-configuration (W-configuration, T-configuration) has as outer face a cycle with $k = 2$ and $c = 1$ ($k = 2$ and $c = 2$, $k = 3$ and $c = 3$), the following corollary follows.

Corollary 1. *Let G be a 1-plane graph that admits an RVR. Then G contains no B-configuration, no W-configuration, and no T-configuration.*

3.4 Proof of sufficient condition

We show that the absence of any B-/W-/T-configuration suffices to compute an RVR. Let G be a 1-plane graph with n vertices and with no B-/W-/T-configuration as a subgraph, see also Fig. 3a for an example. We describe a drawing algorithm, `1PRVDrawer`, which computes an RVR Γ of G. We first assume that G is 3-connected (Sections 3.4.1–3.4.8), and then extend the proof to the general case (Section 3.4.9).

3.4.1 Triangulating and planarizing

The first step of algorithm `1PRVDrawer` is to triangulate G. This is a well-known operation for 3-connected 1-plane graphs, see [1, 17]. However, these algorithms all modify the given 1-planar embedding, even if there is no B-configuration or W-configuration. In fact, this may be required since every crossing must form a kite in a triangulation of G, and so if an edge between endpoints of a crossing already existed, but did not form a face, then it must be re-routed. We do not want to change the embedding, and hence cannot triangulated the 1-plane graph per se. Instead, we combine the triangulation step with a planarization step to ensure that the result has no multiple edges, yet at all crossings (respectively the corresponding dummy-vertices) we have kites.

We construct the triangulated planar graph G^+ with the following steps:

- Let G_1 be the planarization of G. Note that dummy-vertices have degree 4 (and we will maintain this throughout later modifications).
- For each dummy-vertex z of G_1, let $v_0, \ldots, v_3$ be the four neighbors in clockwise order. For $i = 0, \ldots, 3$, consider the face at z between the edges (z, v_i) and (z, v_{i+1}) (addition modulo 4). If this face is the outer face, or an inner face that is not a triangle, then add the edge (v_i, v_{i+1}). Put differently, we add the edges that are needed to turn z into a kite. (In what follows, we will use the previously defined terms "kite", "B-configuration", "W-configuration" and "T-configuration" for the planarized version as well, i.e., for the same situations with crossings replaced by dummy-vertices.) Call the resulting graph G_2. The following is crucial for our correctness:

  Lemma 2. *If G is simple and 3-connected and has no B-configuration, W-configuration or T-configuration, then G_2 has no B-configuration, W-configuration, or T-configuration and no multiple edges.*

  Proof. Every dummy-vertex in G_2 corresponds to a crossing in G. Hence for any W-configuration or T-configuration in G_2, the corresponding crossings in G would define a W-configuration or T-configuration in G since we did not change the embedding or outer face. Now assume that G_2 has a B-configuration, say at dummy-vertex z with neighbors a, b, c, d and with edge (a, b) such that triangle $\{a, b, z\}$ contains c and d inside and hence is not a face. We also surrounded z by a kite, which adds an edge (a, b) for which $\{a, b, z\}$ is a face. Hence there are two copies of edge (a, b).

  It remains to show that there cannot be any double edge (whether created from a B-configuration or otherwise). Assume for contradiction that there is some double edge (u, v).

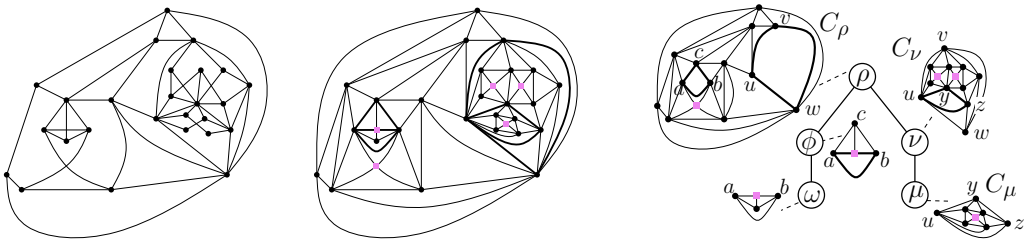

(a) G (b) G^+ (c) $\mathcal{T}$

Figure 3: (a) A 1-plane graph G. (b) The triangulated planar graph G^+ obtained from G. Dummy-vertices are pink squares. (c) The 4-block tree $\mathcal{T}$ of G^+.

It is not possible that the two copies of (u, v) form an inner face, because G was simple, and before adding an edge we always check whether such an edge exists already at this place in the planar embedding. So the two copies either form the outer face or form no face at all. In the latter case, $\{u, v\}$ is a cutting pair of the planar graph G_2, and (as one easily shows) therefore also of G. This contradicts 3-connectivity.

So multiple edges can only exist if the two copies form the outer face. At least one of the copies did not exist in G, and hence was added due to some dummy-vertex z. If the other copy already existed in G, then the crossing at z together with this copy forms a B-configuration in G. If the other copy did not exist in G, then it was added due to some dummy-vertex y, and the crossings at z and y form a W-configuration. □

- In the plane graph G_2 each dummy-vertex is surrounded by a kite by construction. Hence all faces of G_2 that are not triangles contain only original vertices. Triangulate each such face arbitrarily; again this does not create a B-configuration or multiple edges since in a planar graph this would imply a cutting pair.

The resulting graph G^+ is the plane triangulated graph that we use to obtain the RVR. By construction it has the following crucial properties: It has no B-configuration, W-configuration or T-configuration, every dummy-vertex has degree 4, and removing from G^+ all the edges inserted in the above procedure, and replacing each dummy-vertex with a crossing point, we obtain the original 1-plane graph G.

3.4.2 The 4-block tree and an outline

Algorithm `1PRVDrawer` uses a decomposition of G^+ into its 4-connected components. Here, a 4-connected component C of G^+ is a 4-connected triangulated subgraph of G^+. Let the *4-block tree* $\mathcal{T}$ of G^+ be a tree defined as follows (see also [18] and Fig. 3b). Every 4-connected component C_ν of G^+ is represented by a node ν in $\mathcal{T}$. There is an edge between two nodes ν and μ in $\mathcal{T}$, if there is a *separating triangle* (i.e., a triangle with vertices both inside and outside) that belongs to both C_ν and C_μ. We root $\mathcal{T}$ at the node ρ with the 4-connected component that contains the outer face. Then for any parent ν and child μ, the separating triangle common to C_ν and C_μ is an inner face in C_ν and the outer face of C_μ. It is known that $\mathcal{T}$ can be computed in $O(n)$ time [18].

We now give an overview of the remaining steps of algorithm `1PRVDrawer`. It first visits $\mathcal{T}$ top-down and determines for each 4-connected component C_ν a special edge, called *surround-edge* of C_ν. It also computes an RVR γ_ν of C_ν. Next, the algorithm visits $\mathcal{T}$ bottom-up. Let ν be a node of $\mathcal{T}$ with $h \geq 1$ children $\mu_1, \ldots, \mu_h$. Denote by G_ν the graph whose 4-block tree is

the subtree of $\mathcal{T}$ rooted at ν. The already computed visibility representations $\Gamma_{\mu_1}, \ldots, \Gamma_{\mu_h}$ of $G_{\mu_1}, \ldots, G_{\mu_h}$ are suitably merged into γ_ν and an RVR Γ_ν of G_ν is obtained.

At the root we obtain an RVR of $G_\rho = G^+$. In order to turn this in an RVR of G, we must un-planarize, i.e., replace every dummy-vertex by a crossing. For this to be feasible, we need each dummy-vertex is drawn in a special way: its rectangle needs to have exactly one incident edge on each side. We call this the *4-sides-condition*. Even then converting dummy-vertices into crossings is non-trivial, since the four edges on the four sides may not be suitably aligned. To achieve such an alignment, we use the so-called "zig-zag-bend-elimination-slide" [2], with which we can transform parts of the drawing until edges are aligned and hence each dummy-vertex can be turned into a crossing point. Since this introduces no other crossings we obtain an RVR, and deleting the added edges gives the desired RVR of G.

3.4.3 Transversal pairs of bipolar orientations

We will need a few definitions. An *internally 4-connected plane graph* (also known as *irreducible triangulation* [16]) is a plane graph where all inner faces are triangles and that has no separating triangle. Let G be an internally 4-connected planar graph for which the outer face is a 4-cycle. A *transversal pair of bipolar orientations* of G (also known as *regular edge labeling* [19], or simply *transversal structure* [15]) assigns to each inner edge one of two colors[1], say *red* and *blue*, and an orientation such that: (i) For each inner vertex v, the clockwise order of edges around v contains outgoing blue edges, then outgoing red edges, then incoming blue edges, then incoming red edges, and none of these sets is empty. (ii) Let v_N, v_E, v_S, v_W be the four vertices on the outer face of G, in clockwise order. All inner edges incident to $v_N/v_E/v_S/v_W$ are blue incoming/red incoming/blue outgoing/red outgoing. The edges on the outer face can be colored arbitrarily and are oriented so that both the red and the blue graph are acyclic. See e.g. Fig. 4b, where $v_N = v$, $v_E = x$, $v_S = w$, $v_W = u$.

Such a transversal pair of bipolar orientations always exists and can be computed in linear time [16, 19]. The same papers also show how to convert, in linear time, a transversal pair of bipolar orientations into a *rectangular dual representation*. This is a representation of the graph that assigns interior-disjoint rectangles to vertices such that the union of these rectangles is a rectangle without holes. For each edge (v, w), the rectangles of v and w share a positive-length part of their boundaries. Moreover, if the edge is directed $v \to w$, then the rectangle of v is left (below) w if and only if the edge is red (blue). See e.g. Fig. 4c.

3.4.4 Surround-edges

Consider a 4-connected component C_ν with outer face $\{u, v, w\}$ in clockwise order. We want to find an RVR of C_ν. To be able to merge later, it also must satisfy one property: We pick a so-called *surround-edge* e on the outer face beforehand, and we want an RVR such that all edges incident to an endpoint of e are drawn horizontally.

The surround-edge is defined as follows. For any edge e in triangle $\{u, v, w\}$, let x_e be the third vertex of the face of G^+ that is incident to e and inside the triangle. Vertex x_e may be original or a dummy-vertex. We claim that not all three of x_{uv}, x_{vw} and x_{wu} can be dummy-vertices. Assume for contradiction that they are. If they are all distinct, then they would form a T-configuration. If exactly two of them coincide, (say) $x_{uv} = x_{vw}$ is distinct from x_{uw}, then in G the crossing at x_{uw} would form a B-configuration with the edge (u, w) involved in the crossing $x_{uv} = x_{vw}$. Not all three can coincide, else $x_{uv} = x_{vw} = x_{wu}$ would be the only vertex in the graph rooted at ν and have degree 3 and not be a dummy-vertex. So there exists at least one edge e of triangle

[1] In all figures of the paper red edges appear in lighter gray than blue edges when printed b/w.

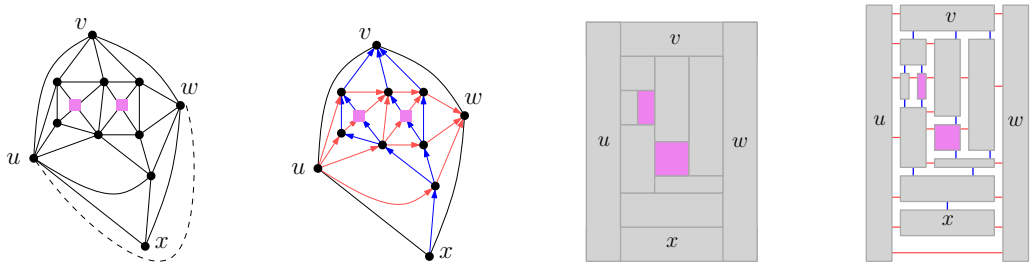

(a) C_ν^- (b) transv.pair (c) $\mathcal{R}$ (d) $\mathcal{R}_\nu$

Figure 4: (a) The graph C_ν^- obtained from C_ν in Fig. 3c. (b) A transversal pair of bipolar orientations for C_ν^-. (c) The corresponding rectangular dual representation. (d) Retracting the rectangles and adding the surround-edge.

$\{u, v, w\}$ where x_e is an original vertex. Define the surround-edge of C_ν to be this edge e. For future reference, we note:

Observation 2. *Let e be the surround-edge of C_ν, and let f_e be the face of G^+ incident to e that is inside the triangle formed by the outer face of C_ν. Then the vertex of f_e that is not an end of e is an original vertex.*

If more than one edge of triangle $\{u, v, w\}$ satisfies the condition of Observation 2, then we chose the surround-edge arbitrarily with one exception: If ν has a parent π, and the surround-edge of C_π is one of the edges on $\{u, v, w\}$, then we use the same edge also as surround-edge of C_ν. Note that whether an edge can be used as surround-edge depends on its incident face in G^+ (not C_π), and so the surround-edge for C_π, if it exists in C_ν, also qualifies as surround-edge for C_ν. For example, for graph C_ϕ in Fig. 3c, edge (a, b) must be chosen as surround-edge, because both (b, c) and (a, c) have a dummy-vertex as third vertex on the face. Since edge (a, b) also belongs to C_ω, it becomes the surround-edge of C_ω as well.

3.4.5 Drawing a 4-connected component

Now we explain how to find an RVR of C_ν such that the surround-edge e and all edges incident to its endpoints are drawn horizontally. If C_ν is K_4, then such an RVR is easily obtained (see also Fig. 7b), so we assume that C_ν has at least 5 vertices. This implies that any inner vertex of C_ν has degree at least 4, else its neighbors would form a separating triangle. Remove the surround-edge from C_ν to obtain graph C_ν^-. Since C_ν is 4-connected, C_ν^- is internally 4-connected. We now obtain an RVR γ_ν of C_ν using a rectangular dual representation as follows:

Lemma 3. *C'_ν has an RVR such that any inner vertex of degree 4 satisfies the 4-sides-condition and all edges incident to an endpoint of the surround-edge of C_ν are drawn horizontally.*

Proof. Since the graph is internally 4-connected, it admits a transversal pair of bipolar orientations, and with it, a rectangular dual representation $\mathcal{R}$ where blue/red edges correspond to shared vertical/horizontal sides [16, 19]. We choose v_W and v_E to be the ends of the surround-edge, and color the outer edges red. Fig. 4b shows the transversal pair of bipolar orientations of the planarization of the graph in Fig. 4a, and Fig. 4c gives the corresponding rectangular dual representation.

Now retract the rectangles a bit as follows. Let $\varepsilon > 0$ be so small that all rectangles have width and height at least ε and the length of any shared side is at least ε. Replace each rectangle

$[x, x'] \times [y, y']$ by a slightly smaller rectangle $[x + \frac{\varepsilon}{3}, x' - \frac{\varepsilon}{3}] \times [y + \frac{\varepsilon}{3}, y' - \frac{\varepsilon}{3}]$. By choice of ε each rectangle still has positive width and height. For each edge of G, the shared side of length at least ε is replaced by a region between its two endpoints, not containing other vertices, that has dimension $2\frac{\varepsilon}{3} \times \frac{\varepsilon}{3}$ or more. Hence for each edge we obtain a line of sight between the rectangles of its endpoints, and hence an RVR with thick lines of sights. Note that the line of sight is horizontal if and only if the shared rectangle-sides were vertical, so the horizontal/vertical lines of sights correspond to the red/blue edges of the transversal pair of bipolar orientations. See also Fig. 4d.

It remains to prove the claim on the 4-sides-condition. Let z be an inner vertex of degree 4, and let R_z be the rectangle of z in the rectangular dual. By properties of the transversal pair of bipolar orientations, the clockwise order of edges around z consists of incoming red edges, incoming blue edges, outgoing red edges, and outgoing blue edges. None of these sets is empty for an inner vertex, and by $\deg(z) = 4$ hence there exists exactly one edge of each kind. In consequence, in the rectangular dual representation there is exactly one neighbor for each side of R_z, which shows the 4-sides-condition. □

This gives an RVR of C_ν^-, to which we need to add the surround-edge. Say the outer face of C_ν^- was u, v, w, x, with surround-edge (u, w) and x the vertex incident to the inner face at the surround-edge. The construction draws all edges at u and w horizontally. After rotation, assume that u is left of w. We can now extend both u and w downwards beyond the drawing, and insert here a horizontal segment for the surround-edge. This gives an RVR $\mathcal{R}_\nu$ of C_ν. See also Fig. 4d.

Observation 3. *In the RVR of C_ν, the 4-sides-condition holds for every dummy-vertex that is an inner vertex of C_ν and has degree 4 in C_ν.*

Proof. The claim holds vacuously if C_ν is K_4, so assume it is not. Let z be a dummy-vertex that is an inner vertex of C_ν. We know that $\deg(z) = 4$ in G^+. This implies that $\deg(z) \leq 4$ in C_ν, and as argued above $\deg(z) \geq 4$ since $C_\nu \neq K_4$. Deleting the special edge does not change the degree of z, therefore $\deg(z) = 4$ even in C'_ν. So if z is an inner vertex of C_ν^-, then the construction for C_ν^- ensures the 4-sides-condition. The only vertex that is an inner vertex of C_ν, but not of C_ν^-, is the vertex x incident to the inner face at the surround-edge. However, by Observation 2 vertex x is not a dummy-vertex, so $z \neq x$. □

3.4.6 Zig-zag-slides

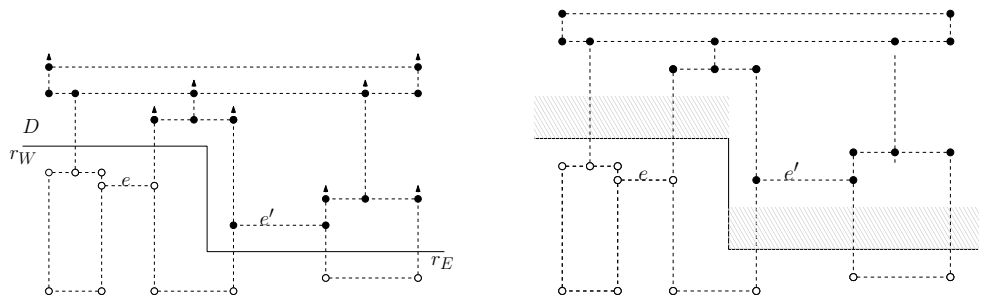

Figure 5: A zig-zag-slide (adapted from [2]). Black points move upward while white points remain stationary. With this amount of sliding, edges e and e' become aligned.

Both for the merging step and for un-doing the planarization, we need a method of modifying the drawing such that some items are moved while others are stationary. This can be achieved with what was called zig-zag-bend-elimination-slide in [2] (we call it a *zig-zag-slide* for short). We

briefly review this here. Assume that Γ is an RVR and we have a dividing curve D as follows: D consists of some vertical line segment s that intersects no horizontal element of the drawing (i.e., neither a horizontal edge nor a horizontal boundary of a rectangle). At the top end of s, attach a leftward horizontal ray r_W. At the bottom end of s, attach a rightward horizontal ray r_E. No conditions are being put onto r_W and r_E. Define the *region above* D be all points that are in the x-range of r_W and strictly above r_W, and all points in the x-range of r_E and on or above r_E.

Now slide all points in the region above D upwards by some amount $\delta > 0$. This maintains a visibility representation, after lengthening vertical edges and rectangle borders suitable, because any horizontal segment either stays stationary or moves in its entirety (recall that no horizontal segments cross s). In particular, for any point p above r_E and any point q above r_W but with a larger y-coordinate, a shift by $y(q) - y(p)$ achieves that these points become horizontally aligned. See Fig. 5. We will use such zig-zag-slides for this shape of D as well as for the three other shapes achieved by flipping the picture horizontally and/or rotating 90°.

3.4.7 Merging components

Let μ be a child of ν in $\mathcal{T}$, and recall that we already obtained drawings γ_ν of C_ν and (recursively) γ_μ of the graph G_μ consisting of C_μ and the components at its descendants. The common vertices $\{u, y, z\}$ of G_μ and C_ν form the outer face of G_μ. See Fig. 3c. The outer face of C_μ^-, consists of $\{u, y, z\}$ as well as a fourth vertex, say x'; after possible renaming of $\{u, y, z\}$ we may assume that the order around the outer face of C_μ^- is u, y, x', z. In particular, drawing γ_μ contains (up to rotation) node u on the top, node y on the right, the surround-edge (y, z) and node x' at the bottom, and node z on the left. Let γ'_μ be the drawing obtained from γ_μ by deleting u, y, z. It suffices to merge γ'_μ, since u, y, z and the edges between them are represented in γ_ν already. See Fig. 6a.

Triangle $\{u, y, z\}$ is an inner face of C_ν. The challenge is now to find a region inside where this face is drawn in γ_ν into which γ'_μ will "fit". Put differently, we want to find a region within γ_ν inside the face $\{u, y, z\}$ that is (after possible rotation) below u, to the left of y and to the right of z. We call such a region, a *feasible region* for γ_μ in γ_ν. As we have no control over how triangle $\{u, y, z\}$ is drawn in γ_ν, the existence of a feasible region is non-trivial, and may in fact require modifying γ_ν slightly. We start with a simple case:

Lemma 4. *Assume that the surround-edge of C_ν also belongs to C_μ. Then there exists a feasible region for γ_μ in γ_ν.*

Proof. If the surround-edge e of C_ν belongs to C_μ, then by our choice of surround-edge we also used e as the surround-edge for C_ν. That is, (y, z) is the surround-edge for both C_μ and C_ν. In γ_ν, (y, z) was drawn (after possible rotation) bottommost, with u above it. A feasible region is now found in the small strip above the drawing of edge (y, z). See Fig. 6b. □

We now show how to find a feasible region for γ_μ in γ_ν.

Lemma 5. *A feasible region for γ_μ in γ_ν always exists, possibly after modifying γ_ν without changing angles or incidences.*

Proof. We already argued this for the case where $\{u, y, z\}$ includes the surround-edge e of C_ν, so assume that this is not the case. Then face $\{u, y, z\}$ of C_ν is also a face of $C'_\nu = C_\nu - e$. The drawing of C'_ν was obtained via the transversal pair of bipolar orientations of C'_ν, with red/blue edges corresponding to horizontal/vertical edges. It is well-known that any inner face in a transversal pair of bipolar orientations has at least one red and at least one blue edge. So $\{u, y, z\}$ is drawn with at least one horizontal and at least one vertical edge, say two edges are

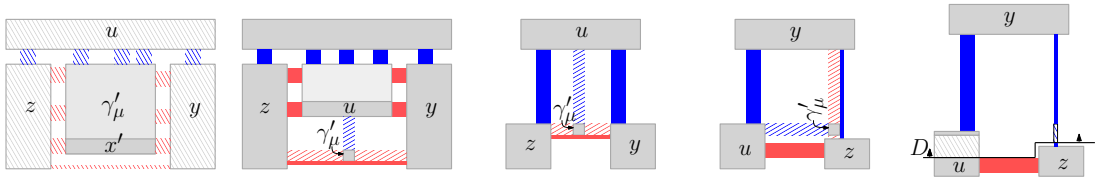

(a) (b) (c) (d) (e)

Figure 6: (a) Illustration of γ_μ and γ'_μ. (b) Illustration for the proof of Lemma 4. (c–e) Illustration for the proof of Lemma 5. (c) u has two vertical edges. (d) y has two vertical edges, u has a higher top than z. (e) Slide so that u has a higher top than z.

vertical and one is horizontal (the other case is similar). Now consider which vertex of $\{u, y, z\}$ is the one where both incident edges of triangle $\{u, y, z\}$ are vertical.

If vertex u has two incident vertical edge (u, y) and (u, z), then (y, z) is drawn horizontally. We can then find a feasible region within the thick line of sight of edge (y, z). See Fig 6c.

Now assume that y has two incident vertical edge (y, u) and (y, z) (the case of vertex z is symmetric). For ease of description, assume that y is above u and z, and u is to the left of z. If the top side of u is strictly above the top side of z, then it is easy to find a feasible region within the thick line of sight of edge (y, z) into which we can merge γ'_ν after rotating it; see Fig. 6d. If the top side of u is below (or at the same height) of the top side of z, then we modify the drawing to create a suitable region using a zig-zag-slide. The dividing curve D is defined as follows. Insert a vertical edge just left of z, ranging from just above (u, z) (hence below the top of u since we have thick lines of visibility) to just above the top of z. Expand it with a horizontal ray leftward from the lower end and a horizontal ray rightward from the upper end. By sliding upward sufficiently far, we hence achieve that the top of u is above the one of z, which allows us to find a feasible region as previously. See Fig 6e. □

By Lemma 5, for every child μ of ν we can find a feasible region into which we merge (after suitable scaling) the drawing $\gamma_\mu - \{u, y, z\}$. This gives the RVR of $C_\nu \cup G_\mu$. It remains to prove the 4-sides-condition for the inner vertices of $C_\nu \cup G_\mu$.

Observation 4. *In the RVR of $C_\nu \cup G_\mu$, the 4-sides-condition holds for every dummy-vertex that is an inner vertex of $C_\nu \cup G_\mu$.*

Proof. Let v be an inner dummy-vertex of $C_\nu \cup G_\mu$. If v is an inner vertex of either C_ν or G_μ, then the statement follows by Observation 3 (or induction on the height of $\mathcal{T}$). Else, v must be on the outer face of C_μ, and hence be part of a separating triangle $\{u, y, v\}$ which is an inner face of C_ν.

Vertex v must have neighbors both inside and outside triangle $\{u, y, z\}$ by 3-connectivity of the triangulated plane graph G^+. Since dummy-vertices have degree 4, it therefore has exactly one neighbor, say x, in $C_\mu - \{u, y, v\}$ and exactly one neighbor, say z, in $C_\nu - \{u, y, v\}$. Further, these neighbors appear in clockwise order z, u, x, y around v and form a kite. See Fig. 7a. It follows that $\{u, y, z\}$ is also a separating triangle (or the outer face of G^+) and that $\{v, z, u\}$ and $\{v, y, z\}$ are faces due to the kite. Therefore C_ν is a K_4, with μ the only child of ν. Similarly, $\{u, x, y\}$ is either a separating triangle or an inner face of G^+, and C_ν is also a K_4 with at most one child (call it ω if it exists).

When choosing the surround-edge of C_ν, we cannot use (u, z) or (y, z), because the inner faces at them contain the dummy-vertex v. Therefore the surround-edge of C_ν is (y, u), and by our choice of surround-edge therefore (y, u) also becomes the surround-edge of C_μ and (if it exists)

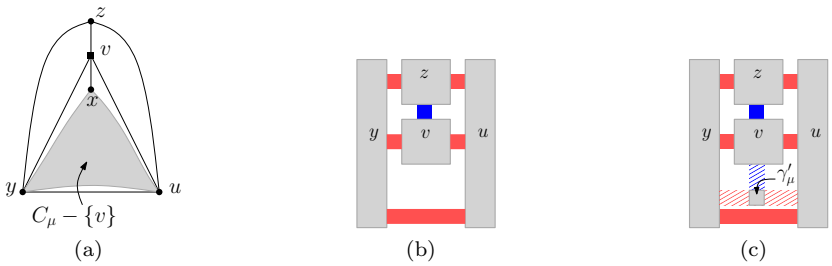

Figure 7: Illustration for the proof of Observation 4.

C_ω. So Lemma 4 is used to find the feasible region for merging G_μ into C_ν. Also, the drawing of C_ν is fixed since it is K_4. Inspecting the construction (see also Fig. 7c) shows that during this merge v obtains an incident edge on each side as desired. □

After repeating the process for the entire 4-block tree we obtain the desired RVR Γ^+ of G^+. After every merging step inner dummy-vertices of degree 4 satisfy the 4-sides-condition. Once we have merged all subgraphs, all inner dummy-vertices have degree 4, and we get the following result:

Observation 5. *In the RVR of G^+, the 4-sides-condition holds for every dummy-vertex that is an inner vertex of G^+.*

3.4.8 Undoing the planarization

Once all merging is completed, the drawing Γ_ρ at the root gives an RVR of G^+. To turn this into an RVR of G, we must undo the planarization, hence remove the dummy-vertices and replace them with a crossing. Since dummy-vertices are surrounded by a kite, none of them is on the outer face of G^+. So any dummy-vertex z is an inner vertex of G^+ and by Observation 5 its rectangle R_z satisfies the 4-sides-condition. Let e_W, e_N, e_E and e_S be the four edges incident at the west/north/east and south side of R_z. If e_W and e_E have the same y-coordinate, and e_N and e_S have the same x-coordinate, then we can simply remove R_z, extend the edges, and obtain the desired crossing.

So assume that e_W and e_E have different y-coordinates, with (say) the y-coordinate y_W of e_W larger than the y-coordinate y_E of e_E. Construct a dividing curve D by starting with a vertical line segment inside R_z, ranging from y-coordinate $y_W + \varepsilon$ to y-coordinate $y_E - \varepsilon$, where $\varepsilon > 0$ is chosen small enough that the segment is inside R_z. Attach a leftward horizontal ray at the top and a rightward horizontal ray at the bottom. Now apply a zig-zag-slide of length $y_W - y_E$. This moves e_E upward while keeping e_W stationary, and hence aligns the two edges. See also Fig. 5, where this is illustrated for $e_W = e$ and $e_E = e'$. With another zig-zag-slide (with the dividing curve using a horizontal line segment and vertical rays), we similarly can align e_N and e_S, if needed. After this, remove the dummy-vertex to obtain the crossing. Repeating this at all dummy-vertices, and finally deleting all added edges, gives the RVR of G, as no crossings are created except where dummy-vertices were removed. This ends the description of algorithm `1PRVDrawer`.

3.4.9 Graphs that are not 3-connected

We now explain the (minor) modifications needed to make algorithm `1PRVDrawer` work even for a 1-plane graph G that is not 3-connected. We used 3-connectivity only in Section 3.4.1 where

we planarized G and then triangulated the result to obtain G^+. If G is 3-connected, then G^+ is a simple graph. If G is not 3-connected, then G^+ may have multiple edges. We now explain how to remove these multiple edges by flipping some of them suitably in such a way that the resulting graph still has an RVR that satisfies the 4-sides-condition, although some dummy-vertex may have degree greater than four.

So assume (v, w) is a multiple edge in the planar graph G^+. At most one copy of (v, w) may have existed in G as well; we call this the *original copy* and all other copies *added copies*. For any added copy e', let the two vertices *facing* e' be the two vertices y, z that are not endpoints of e' but are on a face incident to e'.

If for any added copy of (v, w) both vertices facing it are original, then we can simply *flip* the edge, i.e., remove it and connect these two vertices instead. This does not create a new multiple edge. In particular therefore we may assume that for any added copy at least one facing vertex is a dummy-vertex.

Enumerate the copies of (v, w) as $e_1, \ldots, e_s$ in clockwise order around v in such a way that the 2-cycle e_1, e_s encloses all other copies. We say that e_i has a *dummy on the left/right side* if the vertex that faces e_i and is before/after e_i in the order around v is a dummy-vertex. Observe that if e_i has a dummy on the left side, then no e_j for $j < i$ can have a dummy on the right side, else the two corresponding crossings would form a W-configuration. Also, e_j cannot be the original copy, else it and the crossing at e_i would form a B-configuration. So if ℓ is the maximal index where e_ℓ has a dummy on the left side, then $e_1, \ldots, e_{\ell-1}$ all are added copies and have their crossing on the left side. (We set $\ell = 0$ if there is no such copy.) Likewise if r is the minimal index such that e_r has a dummy on the right side, then $e_{r+1}, \ldots, e_s$ all are added copies and have their crossing on the right side. (We set $r = s + 1$ if there is no such copy.) See Fig. 8.

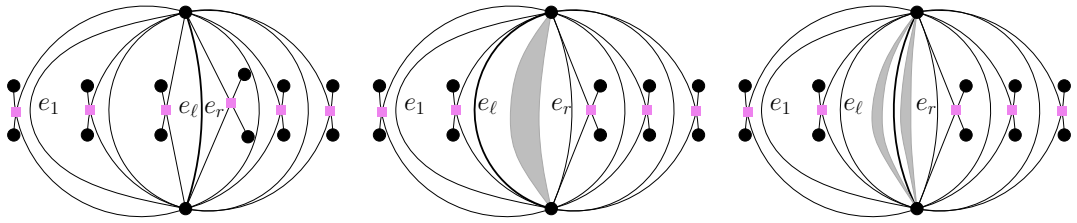

Figure 8: A set of multi-edges. We show $\ell = r$, $\ell = r - 1$ and $\ell = r - 2$. The original edge (or the edge declared to be original) is bold.

Note that $\ell \leq r$ by the above discussion. We consider possibilities for $r - \ell$:

- We may have $\ell = r$. In this case, $e_\ell = e_r$ is the only edge that could be the original copy. If it was not original, then we add e_ℓ to G; note that this does not create any B-configuration due to our careful choice of e_ℓ. We can henceforth assume that the original edge existed.
- We may have $r = \ell + 1$. In this case, one of e_ℓ and e_r may have been the original copy. If neither of them was original, then as in the previous case we add e_ℓ to G.
- We may have $r = \ell + 2$. In this case, edge $e_{\ell+1}$ must have been the original copy, since it has no facing incident dummy-vertex but was not removed by earlier flips.
- We may not have $r > \ell + 2$, else $e_{\ell+1}$ and $e_{\ell+2}$ would both be original, contradicting simplicity of G.

Let G' be the graph obtained by possibly adding e_ℓ to G if we are in one of the first two cases above. We know that G' has no B-configuration, and it has no W-configuration or T-configuration

either since these cannot be created by adding non-crossed edges. The very same graph G^+ also works as planarized triangulated version of G' (the only thing that has changed is that for every multiple edge there now exists an original copy). With this, we now have the following properties:

Observation 6. *Let (u, v) be a multiple edge, and let e be the original copy of (u, v). Then for every other copy e' of (u, v), exactly one of the two vertex y, z facing e' is a dummy-vertex. Moreover, this dummy-vertex forms a separating triangle when combined with e, and edge e' is inside this separating triangle.*

Let $G^{+'}$ be the graph obtained from G^+ by flipping, for every multiple edge, all added copies. This cannot create new multiple edges by planarity. So if we can find an RVR of $G^{+'}$, then this also gives one of the planarization of G after deleting extra edges. However, to undo the planarization we must ensure that the 4-sides-condition holds, and this is non-trivial since flipping an added copy of a multiple edge results in dummy-vertices whose degree exceeds 4.

We claim that for any such dummy-vertex, the 4-sides-condition nevertheless holds automatically when drawing $G^{+'}$ with algorithm `1PRVDrawer`. To see this, let e' be an added copy of an original edge e, let z be the dummy-vertex and y be the original vertex that face e'. When flipping edge e', we insert a new edge (z, y). This edge increases the degree at z, but since y is not a dummy-vertex, this is the only way in which degrees at dummy-vertices can be increased. In particular, (y, z) is inside the separating triangle formed by z and the original edge e and belongs to some child-component C_μ of this separating triangle. The parent-component C_ν contains the dummy-vertex z and the four neighbours at the endpoints of the crossing that defined z, but it does not contain edge (y, z), or any other edge that may have been added at z due to flipping. Therefore when creating the RVR of C_ν, the 4-sides-condition is satisfied for dummy-vertex z, using (parts of) edges that existed in G. Merging child-components may add more edges at z, but no merging step changes edge-directions, and so the four incident edges of z from G remain on the four sides. Thus if we delete all added edges in the RVR of $G^{+'}$, then we obtain an RVR of the planarization of G that satisfies the 4-sides-condition at all dummy-vertices, and can hence undo the planarization exactly as before.

Lemma 6. *Let G be a 1-plane graph G and with no B-/W-/T-configuration as a subgraph. Then algorithm `1PRVDrawer` computes an RVR of G.*

Corollary 1 and Lemma 6 imply Theorem 2. The next two sections prove Theorem 3.

3.5 Area and run-time considerations

Algorithm `1PRVDrawer` does not obviously have linear run-time, since we repeatedly change the entire drawing, especially when applying zig-zag-slides. Also, coordinates may get very small when merging, resulting (after re-scaling to be integral) in a very large area. However, both of these become non-issues if we store RVRs implicitly, using the orthogonal representations of Observation 1. Using this approach, no zig-zag-slide needs to be performed, since these only change edge lengths (but not angles or adjacencies) and hence give rise to exactly the same orthogonal representation. Therefore, all that is required to merge a subgraph at a separating triangle $\{u, y, z\}$ is to find, for each of $w \in \{u, y, z\}$, the appropriate segment along the boundary of the rectangle of w in C_μ, and to merge here the list of adjacencies and angles that w has in the PSLG of C_ν. This takes constant time for w, hence constant time for merging at one separating triangle, and hence linear time overall. All other parts of the algorithm (such as finding separating triangles, computing 4-connected components, finding a transversal pair of bipolar orientations, and finding the rectangular dual representation) take linear time as well.

The final conversion of the PSLG into an RVR also takes linear time, and gives linear integral coordinates. Here "linear" measures the size of the PSLG, which in our case is proportional to the number of vertices, edges, and crossings. For 1-planar graphs this is $O(n)$ since 1-planar graphs have at most $4n - 8$ edges and at most $n - 2$ crossings [8]. Therefore the final RVR has linear coordinates and the area is $O(n^2)$. We conclude:

Lemma 7. *Let G be a 1-plane graph G with n vertices and with no B-/W-/T-configuration as a subgraph. Then there exists an $O(n)$-time algorithm that computes an RVR Γ of G that has $O(n^2)$ area.*

3.6 Testing Algorithm

In order to prove Theorem 3, we need to show how to test whether a given 1-plane graph G contains any B-configuration, any W-configuration, or any T-configuration. Afterwards, if G contains none of them, we can apply algorithm `1PRVDrawer` to produce the desired representation. Hong *et al.* [17] show an algorithm to detect every possible B-configuration and W-configuration in $O(n)$ time, where n is the number of vertices of G. Thus, in what follows we describe how to detect a T-configuration, assuming that G contains no B-configuration and no W-configuration.

Let G^+ be the plane triangulated graph obtained from G by applying the planarization and triangulation step described in Section 3.4. Recall that G^+ contains a T-configuration if and only if G does. Suppose G^+ contains a T-configuration t with outer vertices a, b, c and inner vertices d, e, f, g, h, i, as in Fig. 2d. Since we added a kite around each crossing, vertices a, b, c induce a triangle in G^+. Thus to detect a T-configuration in G^+, we first list all triangles of G^+. The triangle is a T-configuration if and only the three faces on the inside have dummy-vertices as their third vertex, which can be tested in $O(1)$ time per triangle. The only difficulty is hence to list all triangles efficiently. If G^+ is simple, then this can be done in $O(n)$ time (see e.g. [6]). But if G is not 3-connected, then G^+ need not be simple and in particular need not have constant arboricity (the vital ingredient in [6]). Consider again a set of multi-edges as in Figure 8. Observe that if any crossing of an added copy was part of a T-configuration, then so is either the crossing at e_ℓ or the crossing at e_r. Hence to detect a T-configuration, we only need to keep the original copy and the added copies e_ℓ and e_r before computing all triangles. We hence need to find all triangles in a planar graph where edges have multiplicity at most 3; this graph has a linear number of edges in any induced subgraph and hence constant arboricity, and so with [6] finding the triangles can be done in linear time.

Lemma 8. *Let G be a 1-plane graph with n vertices. There exists an $O(n)$-time algorithm to test whether G has a T-configuration, and hence an $O(n)$-time algorithm to test whether G admits an RVR.*

Lemma 7 and Lemma 8 imply Theorem 3.

4 Variable Embedding Setting

Motivated by our characterization, in this section we investigate the variable embedding setting. We remark that, differently from planar graphs, a 1-planar graph can have exponentially many different embeddings, even if 3-connected or 4-connected (see, e.g., Fig. 9). Thus, an exhaustive analysis may not be feasible. Nevertheless, one may wonder whether being free to change the embedding of a given 1-plane graph G is enough to remove every forbidden configuration. In fact, Alam *et al.* [1] proved that if G is 3-connected, then it is possible to reroute its edges so as to

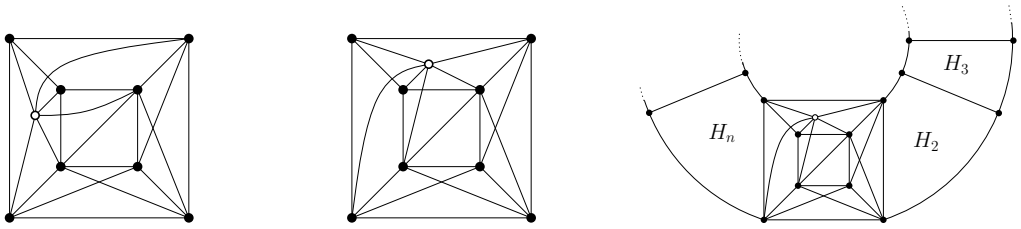

(a) (b) (c)

Figure 9: (a) and (b) show two different 1-planar embedding of the same 4-connected 1-planar graph H. Attaching n copies of H as in (c) we obtain a 4-connected 1-planar graph with $\Theta(n)$ vertices and $\Omega(2^n)$ different 1-planar embeddings.

remove all B-configurations and W-configurations but at most one. We show now that this is not true for T-configuration.

Theorem 4. *Let $k \geq 1$ be an integer value. There exists a 3-connected 1-planar graph G_k with $n = 6k$ vertices that contains at least $k - 1$ T-configurations in every possible 1-planar embedding.*

The following stronger corollary can also be proved.

Corollary 2. *Let $k \geq 1$ be an integer value. There exists a 3-connected 1-planar graph G_k with $n = 6k$ vertices that does not admit any 1-planar embedding that can be represented as an RVR unless we delete at least $k = n/6$ edges.*

The proofs of Theorem 4 and Corollary 2 are postponed to Section 4.1.

On the positive side, a 4-connected 1-planar graph G admits an embedding G' with no B-configurations and W-configurations except that the outer face may consist of a crossing with an edge between two of its ends [1]. In any T-configuration in a 1-plane graph, the three outer vertices are either the outer face or form a separating triplet, so any embedding of a 4-connected 1-planar graph has at most one T-configuration. This means that in G' there is at most one T-configuration, and if there is one then there is no B-configuration or W-configuration. This leads to the following theorem.

Theorem 5. *Every 4-connected 1-planar graph G admits a 1-planar embedding that can be represented as an RVR, except for at most one edge.*

An *optimal 1-planar graph* is one that has the maximum number $4n - 8$ of possible edges. Optimal 1-planar graphs are always 4-connected, and in any 1-planar embedding all crossings are in kites, with the exception of one crossing on the outer-face [27]. It follows that G does not contain any T-configuration, and with the exception of the crossing at the outer face it contains no B-configuration or W-configuration.

Corollary 3. *Every optimal 1-plane graph G admits an RVR, except for one edge.*

4.1 Proof of Theorem 4 and Corollary 2

The idea is to show that K_6 has a "unique" embedding, up to renaming and choice of the outer face. To argue this, define the *planar skeleton* $S(G)$ of a 1-plane graph G to be the plane graph obtained from G by removing all edges that cross some other edge.

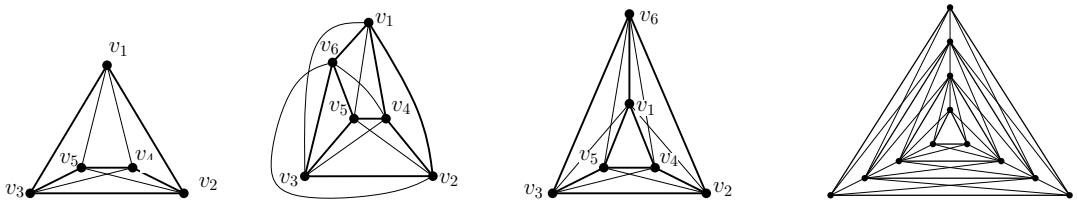

(a) K_5 (b) K_6 (c) K_6 (d) G_k

Figure 10: (a) K_5 in its unique (up to renaming and outer face) 1-planar embedding. (b-c) K_6 in its unique (up to renaming) 1-planar embedding and two possible choices of the outer face. (C) Multiple copies of K_6 combined.

Observation 7. *For any 1-planar embedding of K_6, the planar skeleton is the 3-prism and the outer face forms a B-configuration or a T-configuration.*

Proof. Recall that the 3-prism is a 3-connected planar graph whose unique planar embedding (up to renaming and outer face) is shown in Fig. 10b (the bold edges only).

The complete graph K_5 has a unique 1-planar embedding (up to renaming and outer face) [21], which is shown in Fig. 10a. Observe that every face of such embedding is a triangle, either incident to two vertices and a crossing point or to three vertices. In order to realize a 1-planar embedding of K_6, a sixth vertex can only be added to one of the four faces of K_5 that are incident to three vertices. Two of these four cases are illustrated in Figs. 10b and 10c; the other two cases are symmetric. In every one of these two cases the statement holds. □

Corollary 4. *K_6 has no 1-planar embedding that can be represented an an RVR.*

Proof. Consider an arbitrary 1-planar embedding of K_6. Since the planar skeleton of K_6 is a 3-prism, every planar edge is incident to a face of degree 4 in the 3-prism, and all such faces contain a crossing in K_6. If the outer face contains a crossing then this gives a B-configuration (Fig. 10b), and if it consists of vertices only then it gives a T-configuration (Fig. 10c). Either way it does not admit an RVR by Lemma 1. □

Now for any k we can create a 3-connected 1-planar graph G_k with $6k$ vertices that consists of k vertex-disjoint copies of K_6. Fig. 10d shows one possible such construction for $k = 2$. Also, it is immediate to see that all copies of K_6 in G_k, except for at most one, must be embedded with no crossing points on their outer face (i.e., as in Fig. 10c). Thus, G_k contains at least $k - 1$ T-configurations. This concludes the proof of Theorem 4.

Moreover, if some spanning subgraph H of G had an RVR Γ, then for each of the k copies of K_6, at least one edge cannot exist in Γ, else Γ would induce an RVR of K_6. Hence at least $k = n/6$ edges of G are not represented in Γ. This proves Corollary 2.

5 Conclusions and Open Problems

In this paper, we studied rectangle visibility representations of non-planar graphs with a fixed embedding. We showed that we can test in polynomial time whether a graph has such a representation. Of special interest to us were 1-planar graphs; here we can give a linear-time algorithm to test the existence of visibility representations if the embedding is fixed. Moreover, in case of a negative answer the algorithm provides a witness in form of either a B-configuration, W-configuration, or T-configuration. We also briefly studied 1-planar graphs without fixed

embeddings; we showed that not all 3-connected 1-planar graphs have visibility representations with at most one crossing per edge, but all 4-connected 1-planar graphs do after deleting one edge.

The most pressing open problem is whether we can restrict the drawings less and still test for the existence of visibility representations? Most importantly, if we fix the rotational scheme and outer face, but *not* order in which crossing occur (and perhaps not even which edges cross), is it possible to test whether a visibility representation respecting the rotational scheme and outer face exists? The NP-hardness proof of Shermer [25] does not hold if the rotational scheme is fixed, since it uses a reduction from linear-arboricity-2, and fixing the rotational scheme would severely restrict the possible ways of splitting a graph into two linear forests. The orthogonal representation approach utterly fails if the order of crossings is not fixed, since it requires planarization as a first step.

Secondly, can we characterize for more graph classes exactly when they have a rectangle visibility representation? Notice that Lemma 1 does not use 1-planarity, and hence gives a necessary condition for any subgraph of a graph G (with a fixed embedding) to have an RVR. A second necessary condition stems from that as soon as edges may have 2 or more crossings, the edge-parts between the crossings may be non-trivial and have cycles; clearly the graph formed by these crossings must be bipartite if G has a visibility representation. Are these two conditions sufficient, and if not, can we find necessary and sufficient conditions, at least for some restricted graph classes such as 2-planar and fan-planar graphs?

Finally, it is immediate to observe that any RVR can be transformed into a drawing such that vertices are points, edges are polylines with at most 2 bends, and crossings occur only at right angles. Hence, Lemma 7 can be used to compute *RAC-drawings* (see e.g. [11]) with at most 2 bends per edge in linear time and quadratic area. Since not all 1-plane (straight-line drawable) graphs admit a bendless RAC-drawing [13] (and when they do they may require exponential area [5]), we ask whether all 1-plane graphs admit a RAC-drawing with at most 2 bends per edge in polynomial area?

References

[1] M. J. Alam, F. J. Brandenburg, and S. G. Kobourov. Straight-line grid drawings of 3-connected 1-planar graphs. In *GD 2013*, volume 8242 of *LNCS*, pages 83–94. Springer, 2013.

[2] T. Biedl, A. Lubiw, M. Petrick, and M. J. Spriggs. Morphing orthogonal planar graph drawings. *ACM Trans. Algorithms*, 9(4):29, 2013.

[3] P. Bose, A. M. Dean, J. P. Hutchinson, and T. C. Shermer. On rectangle visibility graphs. In *GD 1996*, volume 1190 of *LNCS*, pages 25–44. Springer, 1997.

[4] F. J. Brandenburg. 1-Visibility representations of 1-planar graphs. *J. Graph Algorithms Appl.*, 18(3):421–438, 2014.

[5] F. J. Brandenburg, W. Didimo, W. S. Evans, P. Kindermann, G. Liotta, and F. Montecchiani. Recognizing and drawing IC-planar graphs. In *GD 2015*, LNCS. Springer, 2015. In press.

[6] N. Chiba and T. Nishizeki. Arboricity and subgraph listing algorithms. *SIAM J. Comput.*, 14(1):210–223, 1985.

[7] S. Cornelsen and A. Karrenbauer. Accelerated bend minimization. *J. Graph Algorithms Appl.*, 16(3):635–650, 2012.

[8] J. Czap and D. Hudák. On drawings and decompositions of 1-planar graphs. *Electr. J. Comb.*, 20(2):P54, 2013.

[9] A. M. Dean, W. S. Evans, E. Gethner, J. D. Laison, M. A. Safari, and W. T. Trotter. Bar k-visibility graphs. *J. Graph Algorithms Appl.*, 11(1):45–59, 2007.

[10] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice-Hall, 1999.

[11] W. Didimo and G. Liotta. The crossing angle resolution in graph drawing. In J. Pach, editor, *Thirty Essays on Geometric Graph Theory*. Springer, 2012.

[12] P. Duchet, Y. Hamidoune, M. L. Vergnas, and H. Meyniel. Representing a planar graph by vertical lines joining different levels. *Discrete Math.*, 46(3):319 – 321, 1983.

[13] P. Eades and G. Liotta. Right angle crossing graphs and 1-planarity. *Discrete Appl. Math.*, 161(7-8):961–969, 2013.

[14] W. S. Evans, M. Kaufmann, W. Lenhart, T. Mchedlidze, and S. K. Wismath. Bar 1-visibility graphs vs. other nearly planar graphs. *J. Graph Algorithms Appl.*, 18(5):721–739, 2014.

[15] S. Felsner. Rectangle and square representations of planar graphs. In J. Pach, editor, *Thirty Essays on Geometric Graph Theory*, pages 213–248. Springer, 2013.

[16] É. Fusy. Transversal structures on triangulations: A combinatorial study and straight-line drawings. *Discrete Math.*, 309(7):1870–1894, 2009.

[17] S. Hong, P. Eades, G. Liotta, and S. Poon. Fáry's theorem for 1-planar graphs. In *COCOON 2012*, volume 7434 of *LNCS*, pages 335–346. Springer, 2012.

[18] G. Kant. A more compact visibility representation. *Int. J. Comput. Geometry Appl.*, 7(3):197–210, 1997.

[19] G. Kant and X. He. Regular edge labeling of 4-connected plane graphs and its applications in graph drawing problems. *Theor. Comput. Sci.*, 172(1-2):175–193, 1997.

[20] G. Kant, G. Liotta, R. Tamassia, and I. G. Tollis. Area requirement of visibility representations of trees. *Inf. Process. Lett.*, 62(2):81–88, 1997.

[21] V. P. Korzhik. Minimal non-1-planar graphs. *Discrete Math.*, 308(7):1319–1327, 2008.

[22] R. H. J. M. Otten and J. G. V. Wijk. Graph representations in interactive layout design. In *IEEE ISCSS*, pages 914–918. IEEE, 1978.

[23] J. Pach and G. Tóth. Graphs drawn with few crossings per edge. *Combinatorica*, 17(3):427–439, 1997.

[24] P. Rosenstiehl and R. E. Tarjan. Rectilinear planar layouts and bipolar orientations of planar graphs. *Discr. & Comput. Geom.*, 1:343–353, 1986.

[25] T. C. Shermer. On rectangle visibility graphs. III. External visibility and complexity. In *CCCG 1996*, pages 234–239. Carleton University Press, 1996.

[26] I. Streinu and S. Whitesides. Rectangle visibility graphs: Characterization, construction, and compaction. In *STACS 2003*, volume 2607 of *LNCS*, pages 26–37. Springer, 2003.

[27] Y. Suzuki. Re-embeddings of maximum 1-planar graphs. *SIAM J. Discrete Math.*, 24(4):1527–1540, 2010.

[28] R. Tamassia. On embedding a graph in the grid with the minimum number of bends. *SIAM J. Computing*, 16(3):421–444, 1987.

[29] R. Tamassia and I. G. Tollis. A unified approach to visibility representations of planar graphs. *Discr. & Comput. Geom.*, 1(1):321–341, 1986.

[30] C. Thomassen. Plane representations of graphs. In *Progress in Graph Theory*, pages 43–69. AP, 1984.

[31] C. Thomassen. Rectilinear drawings of graphs. *J. Graph Theory*, 12(3):335–341, 1988.

[32] S. K. Wismath. Characterizing bar line-of-sight graphs. In *SoCG 1985*, pages 147–152. ACM, 1985.